arXiv:1512.01947v1 [stat.ML] 7 Dec 2015

LEARNING POPULATION AND SUBJECT-SPECIFIC BRAIN CONNECTIVITY NETWORKS VIA MIXED NEIGHBORHOOD SELECTION

By Ricardo Pio Monti*, Christoforos Anagnostopoulos* and Giovanni Montana*,†

*Imperial College London** *and King's College London*†

In neuroimaging data analysis, Gaussian graphical models are often used to model statistical dependencies across spatially remote brain regions known as functional connectivity. Typically, data is collected across a cohort of subjects and the scientific objectives consist of estimating population and subject-specific graphical models. A third objective that is often overlooked involves quantifying inter-subject variability and thus identifying regions or sub-networks that demonstrate heterogeneity across subjects. Such information is fundamental in order to thoroughly understand the human connectome. We propose *Mixed Neighborhood Selection* in order to simultaneously address the three aforementioned objectives. By recasting covariance selection as a neighborhood selection problem we are able to efficiently learn the topology of each node. We introduce an additional mixed effect component to neighborhood selection in order to simultaneously estimate a graphical model for the population of subjects as well as for each individual subject. The proposed method is validated empirically through a series of simulations and applied to resting state data for healthy subjects taken from the ABIDE consortium.

1. Introduction. At the forefront of neuroscientific research is the study of functional connectivity; defined as the statistical dependencies across spatially remote brain regions [Friston, 1994, 2011]. While traditional neuroimaging studies focused on the roles of specific brain regions, there has recently been a significant shift towards understanding the connectivity across regions [Smith, 2012]. This shift has been partially catalyzed by recent advances in imaging techniques. In particular, the introduction of functional MRI (fMRI) has played a crucial role by providing a non-invasive mechanism through which to obtain whole-brain coverage of neuronal activity [Huettel, Song and McCarthy, 2004, Poldrack, Mumford and Nichols, 2011]. The focus of this work involves estimating functional connectivity networks from fMRI data, however the methodology presented can also be used in conjunction with other imaging modalities.

From a statistical perspective, Gaussian Graphical models (GGMs) are often employed to model functional connectivity [Smith et al., 2011, Varoquaux and Craddock, 2013]. In this manner, undirected connectivity networks can be inferred by studying the conditional independence structures across brain regions [Lauritzen, 1996].

Keywords and phrases: functional connectivity, neuroimaging, graphical models, inter-subject variability

However, there are several caveats of neuroimaging data which must be adequately considered in order to accurately estimate functional connectivity networks, principal among which is the *replicated* nature of imaging experiments. Neuroimaging datasets typically consist of multivariate time series data collected across a cohort of subjects. As a consequence, the goal is often to both learn multiple *replicated* brain connectivity networks across a cohort of subjects as well as understand the properties and characteristics which define the population. Hereafter we refer to such data as *replicated* data.

Two further hallmarks of neuroimaging data are its high dimensional nature as well as its reproducibility. Particularly in the context of connectivity networks, observed patterns across large numbers of brain regions have been shown to be consistent across subjects and scanning sessions [Damoiseaux et al., 2006, Zuo et al., 2010]. In recent years, these properties have motivated novel methodological advances which seek to more reliably estimate subject-specific functional connectivity networks by leveraging information across a population [Varoquaux et al., 2010, Wee et al., 2014]. However, from a neuroscientific perspective several problems remain; chief of which is understanding and quantifying inter-subject variability [Kelly et al., 2012, Mueller et al., 2013].

Within the neuroimaging literature, standard approaches to studying *replicated* data can be summarized in two main avenues of research. The first involves learning a separate GGM for each subject independently. While such an approach allows researches to examine subject-specific hypotheses, it may also prove problematic in the context of neuroimaging data. This is particularly true when studying fMRI data, which is characterized by its high spatial and low temporal resolution. Much of the neuroimaging literature looks to address this issue via the introduction of regularization constraints [Smith et al., 2011, Ryali et al., 2012, Lee et al., 2011, Wehbe et al., 2015], however such methods may still perform poorly if only a limited number of observations are available per subject. Notwithstanding, a practical advantage such a strategy is that variability across a population can be quantified [Varoquaux and Craddock, 2013]. However, the resulting two-step procedure is ultimately limited by the fidelity of the estimated individual networks.

The second approach is to learn a single GGM that is representative of the entire population of brain networks. In addition to greatly facilitating the interpretation of results, this strategy leverages information across subjects (albeit in a naive manner), thereby alleviating issues caused by the high dimensional nature of the data. The most significant deficiency stems from the fact that it fails to adequately model the inter-subject variability present in neuroimaging data [Poldrack, Mumford and Nichols, 2011, Ashby, 2011, Mueller et al., 2013]. Furthermore, the question of understanding variability across the population is often sidelined thereby undermining the interpretability of results [Fallani et al., 2014].

The objective of this work is to reconcile the two popular approaches presented above; thus allowing for accurate network estimation at subject-specific and population levels while also quantifying variability present across a cohort. In recent years significant progress has been made towards these objectives, in particular through the use of novel regularization schemes. For example, the *graphical lasso* penalty [Friedman, Hastie and

Tibshirani, 2008] has been used extensively to recover the sparse support on a subject-specific basis [Varoquaux and Craddock, 2013]. An alternative approach involves sharing information across subjects via employing either *group* or *fused* lasso penalties [Varoquaux et al., 2010, Danaher, Wang and Witten, 2014]. Based upon the notion that resting state networks demonstrate reproducible properties across subjects, such approaches look to alleviate the high-dimensional nature of fMRI data by effectively leveraging data. This recent work has clearly demonstrated the benefits obtained through leveraging information across a cohort of subjects. However, while such approaches constitute significant improvements both from a methodological perspective, the question of quantifying inter-subject variability remains unanswered.

In this work we present a novel methodology through which to partially address these issues. The proposed method, named Mixed Neighborhood Selection (MNS), is based on the neighborhood selection method introduced by Meinshausen and Bühlmann [2006]. By recasting covariance selection as a series of linear regression problems, neighborhood selection methods are able to learn the local network topology of each region. MNS extends neighborhood selection by incorporating an additional random effect component. This component is introduced with the intention of learning both a population topology (captured in the fixed effect terms) as well as subject-specific topology (captured in the random effects) for each node. This serves to directly model inter-subject variability and provides a much richer model of functional connectivity. In particular, the proposed method is able to partition edges according to their reproducibility across the cohort. In doing so, MNS provides an additional layer of information which can be exploited to further understand functional connectivity. Moreover, by effectively differentiating between reproducible edges present across the entire cohort and highly variable edges, the proposed method is able to leverage information in a discriminative manner, leading to more reliable network estimates.

In order to illustrate the capabilities of the proposed method we present a brief motivating example, shown in Figure [1]. We consider a scenario where there is a population consisting of four individuals whose functional connectivity networks share a common structure but also demonstrate some variability. In particular, one edge varies across subjects such that two subjects exhibit the functional connectivity shown in Figure [1a] and the remaining two Figure [1b]; the edge in question (edge A) is shown to vary from positive to negative across groups. In such a scenario, it is of scientific interest both to uncover the correct functional connectivity networks as well as to correctly identify edges which are variable within the population. This is precisely what MNS is capable to achieving. The results are shown in Figure [1c] where the blue lines indicate edges shared across the entire population. The thick gray edges indicate *random effect* edges that demonstrate high variability. Figure [1d] shows the estimated edge coefficients for two edges of interest when estimated using the proposed method and the graphical lasso. We note that in addition to correctly recovering the sparsity structure, the proposed method is able to discriminate edges according to their reproducibility across the cohort. This is in contrast to what could be achieved by studying the networks estimated for each subject independently. This point is demonstrated in Figure [1d] where the es-

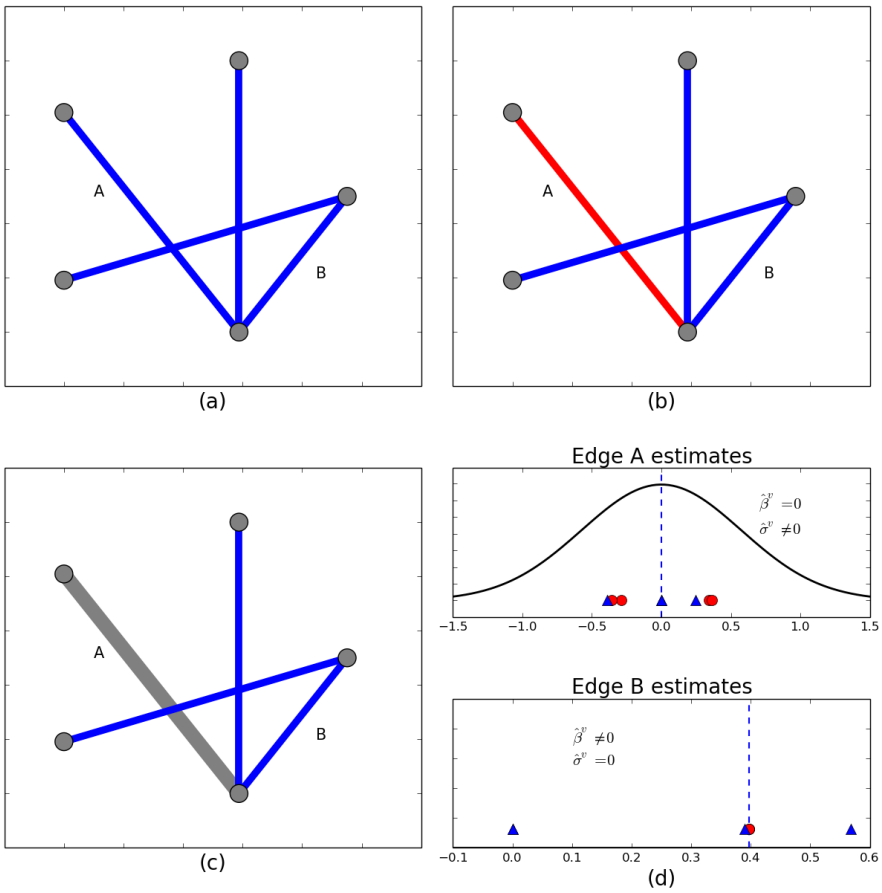

FIG 1. *Toy motivating example to illustrate the capabilities of the proposed method. Networks were simulated with $p = 5$ nodes and with $n = 8$ observations per subject for $N = 4$ subjects. The functional connectivity networks for two of the subjects is shown in* ***(a)*** *while the networks for the remaining two is shown in* ***(b)****. Blue edges indicating positive partial correlations and red edges indicative of negative partial correlations. A significant proportion of the edges are shared across subjects with a single variable edge. The results for our proposed method are shown in* ***(c)****; blue lines indicate edges shared by the entire population while thick gray edges indicate highly variable edges. Estimated edge coefficients for edges A and B are shown as obtained by the MNS algorithm as well as by applying the graphical lasso to each dataset independently in* ***(d)****: Dashed blue lines indicate the estimated population edge value while the solid back line is the estimated probability density function of that edge based on the random effects (i.e., it visualizes the BLUP distribution). Blue, triangular points indicate edge values as estimated by the graphical lasso while red, circular points indicate subject-specific MNS estimates (i.e., the BLUPs).*

timated edge coefficients for the graphical lasso[1] are shown to be highly variable across both edges, one of which is does not vary across subjects. As a result, it follows that

[1] the graphical lasso was run independently for each subject. The regularization parameter for each subject was selected using cross-validation

identifying variable edges in a two-step procedure is challenging, even in low-dimensions.

The remainder of this manuscript is organized as follows. Prior to discussing the proposed method, we briefly review current methodology in neuroimaging in Section 2. The proposed is then detailed in Section 3. In Section 4 we present an extensive simulation study and the proposed method is applied to resting-state fMRI data from the ABIDE consortium in Section 5.

2. Current methodology in Neuroimaging. In this section we outline current methodology within the neuroimaging community. To set notation, we assume we have access to fMRI time series across a cohort of N subjects. For the ith subject, it is assumed we observe an n-dimensional fMRI time series across p fixed regions of interest. We write $V = \{1, \ldots, p\}$ to denote the set of regions or nodes and refer to the dataset for the ith subject by $X^{(i)} \in \mathbb{R}^{n \times p}$. Further, we write $X_v^{(i)} \in \mathbb{R}^{n \times 1}$ to denote the time-series for any node $v \in V$. Similarly, we let $X_{\setminus v}^{(i)} \in \mathbb{R}^{n \times (p-1)}$ denote the times-series across all remaining nodes.

Throughout this work it is assumed that the data of each subject follows a stationary multivariate Gaussian distribution. Since our primary interest is the estimation of functional connectivity networks, summarized in the inverse covariance matrix, we assume without loss of generality that each $X^{(i)}$ has zero mean. As a result, we have that $X^{(i)} \sim \mathcal{N}(0, \Sigma^{(i)})$ where $\Sigma^{(i)}$ is the covariance for the ith subject.

2.1. *Modeling connectivity through GGMs.* Under the assumption of Gaussianity, estimating functional connectivity networks is equivalent to learning the conditional independence structure for each subject. This can be succinctly represented as a graphical model, $G^{(i)} = \left(V, E^{(i)}\right)$, where the edge set, $E^{(i)}$, encodes conditional dependencies across a fixed set of nodes, V [Lauritzen, 1996]. Formally, the edge set summarizes the non-zero entries in the precision matrix, thus:

$$E^{(i)} = supp\left(\left(\Sigma^{(i)}\right)^{-1}\right) = \left\{(j,k) : \left(\Sigma^{(i)}\right)^{-1}_{j,k} \neq 0\right\}. \tag{2.1}$$

The resulting edges are taken to be indicative of functional relationships between spatially remote regions of the brain, allowing us to interpret the estimated graphical model as a functional connectivity network.

The problem of learning the aforementioned dependence structures, first posed by Dempster [1972], is typically referred to as covariance selection. This is a challenging problem that is further exasperated by the high-dimensional nature of fMRI data. As a result, regularization is often introduced [Smith et al., 2011, Varoquaux and Craddock, 2013]. In particular, the use of neighborhood selection methods introduced by Meinshausen and Bühlmann [2006] have been widely adopted within the neuroimaging community [Lee et al., 2011, Wee et al., 2014, Pircalabelu et al., 2015]. As these methods will form the backbone for the proposed method, we formally discuss neighborhood selection below.

2.1.1. *Neighborhood selection.* The intuition behind neighborhood selection stems from the fact that we may learn the conditional independence structure across all nodes by iteratively learning the independence structure of each node. The latter is referred to as the neighborhood for each node $v \in V$. We write $ne^{(i)}(v)$ to denote the estimated neighborhood of node v at the ith subject.

Meinshausen and Bühlmann [2006] propose to learn the neighborhood of each node $v \in V$ by considering the optimal prediction of $X_v^{(i)}$ given the time series of the remaining nodes. This results in the formulation of the following linear model for node v:

$$X_v^{(i)} = X_{\setminus v}^{(i)} \beta^{(i),v} + \epsilon^{(i),v} \tag{2.2}$$

where $\epsilon^{(i),v} \sim \mathcal{N}(0, \sigma^2 I)$ is white noise. In such a regression model it follows that nodes that are not in the neighborhood of v will be omitted from the set of optimal predictors. Thus neighborhood selection can be reformulated as subset selection in a linear model. The latter problem has received considerable attention; one notable solution being that of the Lasso [Tibshirani, 1996]. Briefly, the Lasso imposes a constraint on the ℓ_1 norm of the regression coefficients. This allows the Lasso to obtain parsimonious solutions whilst remaining convex.

The neighborhood selection approach described in Meinshausen and Bühlmann [2006] proceeds by solving the following convex optimization problem for each node v:

$$\hat{\beta}^{(i),v} = \underset{\beta^{(i),v} \in \mathbb{R}^{p-1}}{\operatorname{argmin}} \left\{ \frac{1}{2} ||X_v^{(i)} - X_{\setminus v}^{(i)} \beta^{(i),v}||_2^2 + \lambda ||\beta^{(i),v}||_1 \right\}. \tag{2.3}$$

Due to the parsimony property of the Lasso, some elements of $\hat{\beta}^{(i),v} \in \mathbb{R}^{p-1}$ will be shrunk to zero, effectively removing these nodes from the optimal prediction set. An estimate for the neighborhood of v is subsequently defined as:

$$\hat{ne}^{(i)}(v) = \left\{ u \in V \setminus \{v\} : \hat{\beta}_u^{(i),v} \neq 0 \right\}. \tag{2.4}$$

That is to say, the neighborhood of v is the set of all nodes included in the Lasso solution to equation (2.3). Given an estimate of the neighborhood of all nodes, the edge structure for a graphical model can then be obtained using one of the following rules [Meinshausen and Bühlmann, 2006]:

$$E_{OR}^{(i)} = \{(v,u) : u \in \hat{\beta}^{(i),v} \textbf{ or } v \in \hat{\beta}^{(i),u}\} \text{ or } E_{AND}^{(i)} = \{(v,u) : u \in \hat{\beta}^{(i),v} \textbf{ and } v \in \hat{\beta}^{(i),u}\}. \tag{2.5}$$

2.1.2. *Persistent challenges.* A fundamental challenge associated with estimating functional connectivity networks is that the accuracy of the estimated graphical models remains limited when the number of observations, n, is small relative to the dimensionality of the data, p. This is particularly concerning in the context of fMRI data which is renown for its high spatial and low temporal resolution [Poldrack, Mumford and Nichols, 2011].

While the introduction of regularization is able to partially address this concern, such methods will still perform poorly in the presence of limited observations (see for example the results of our motivating example presented in Figure [1]). Moreover, there is

no consensus within the neuroimaging community regarding the *true* sparsity level of functional connectivity networks. Some studies suggest that connectivity networks have evolved to achieve high efficiency of information transfer at a low connection cost, leading to sparse networks [Bullmore, 2009]. However, convincing arguments have also been posited against sparsity. For example, Markov et al. [2013] propose a high-density model where efficiency is achieved via the presence of highly heterogeneous edge strengths between nodes. As such, it is unclear how sparsity inducing regularization techniques serve the inferential goals of cognitive neuroscience [Wehbe et al., 2015]. This is particularly true in the context of functional connectivity.

A naive solution to this concern is to concatenate observations across all subjects and estimate a single graphical model. While such approaches may lead to interpretable results, unless inter-subject heterogeneity is adequately modeled the resulting network will be a poor representation of the true population covariance structure [Ashby, 2011].

Recently, more sophisticated methods have sought to address this issue by exploiting the fact that functional connectivity networks are often reproducible across subjects [Varoquaux et al., 2010, Guo et al., 2011, Danaher, Wang and Witten, 2014]; thereby suggesting that improved connectivity estimates could be obtained via leveraging information across a cohort of subjects. Such methods effectively regularize the differences between estimated networks across multiple subjects via the introduction of an additional penalty. For example, by imposing a group lasso penalty Varoquaux et al. [2010] are able to enforce a common support structure on the graphical models of each subject.

However, as we look to demonstrate in the remainder of this work, these approaches suffer two fundamental deficiencies. First and foremost, they do not quantify variability across a cohort of subjects. This is a topic which has not received sufficient attention within the neuroimaging community and is fundamental for the understanding and interpretation of estimated networks [Fallani et al., 2014]. Second, the aforementioned methods leverage information across subjects in a homogeneous fashion. As we will demonstrate, the proposed method is able to discriminate edges according to their reproducibility. This information then serves to guide the proposed method when leveraging information; allowing for less stringent regularization to be enforced upon variable edges.

3. Methods. There are a number of challenges and objectives which must be considered when studying neuroimaging data across a cohort of subjects. Predominant objectives include learning population and subject-specific connectivity networks. In this section we present the proposed method, termed Mixed Neighborhood Selection (MNS), through which to address both these objectives together with a third that is often overlooked; that of understanding the variability in covariance structure across a population.

We look to address these objectives by adequately modeling inter-subject heterogeneity. This is achieved via proposing a novel model for the covariance structure across a cohort of subjects. The proposed model, which looks to categorize edges on the basis of their reproducibility, is described in Section 3.1. This in turn allows MNS to leverage information across subjects in a discriminative manner, described in Section 3.2. An efficient estimation algorithm is outlined in Section 3.3. Parameter tuning is dis-

cussed in Section 3.4 and Section 3.5 contains a discussion of related approaches from the literature.

3.1. *A novel covariance model.* We propose to model the covariance structure for each subject as the union of a shared covariance structure together with subject-specific idiosyncrasies. This corresponds to the assumption that there exists a shared covariance structure which manifests itself across all subjects together with subject-specific deviations from this structure. The latter allows our model to accommodate inter-subject variability which cannot be ignored. As a result, we model the conditional independence structure of each subject as the union of the support of a sparse population network and a subject-specific network. Formally, the support for each subject's conditional independence structure, originally defined in equation (2.1), is modeled as:

$$E^{(i)} = E^{pop} \cup \tilde{E}^{(i)} \tag{3.1}$$

Here we interpret E^{pop} as the population edges which encode the conditional independence structure shared across the entire population. Under the assumption of Gaussianity, it follows that E^{pop} is associated with a population precision matrix, $\Theta^{pop} \in \mathbb{R}^{p \times p}$. On the other hand, it is $\tilde{E}^{(i)}$ which encodes subject-specific deviations from the population covariance structure. We define $\tilde{E} = \bigcup_{i=1}^{N} \tilde{E}^{(i)}$ as the set of edges demonstrating variability across the entire population. This variability may either be attributed to the nature of the edge (i.e., positive or negative partial correlations as in the motivating example described in Figure [1]) or partial presence of the edge (i.e., the edge is only present within a sub-group of subjects).

The objective of the proposed method therefore corresponds to accurately identifying both E^{pop} and $\tilde{E}^{(i)}$. Given E^{pop} and $\tilde{E}^{(i)}$, one can infer $E^{(i)}$ and $\tilde{E}$. However, by focusing on E^{pop} and $\tilde{E}^{(i)}$, as opposed to directly considering subject-specific edges, a far richer description of functional architecture is obtained. In the case of the motivating example presented in Section 1, $\tilde{E} = \tilde{E}^{(i)} = \{A\}$ for all subjects while the remaining edges are captured in E^{pop}. From the perspective of neuroimaging, partitioning edges in this manner is fundamental to further understanding the functional architecture of the brain [Kanai and Rees, 2011, Zilles and Amunts, 2013].

It is useful to note that the model described in equation (3.1) generalizes two typical approaches in the study of functional connectivity. The traditional method of estimating a single population network, Θ^{pop}, by concatenating data across all subjects is equivalent to the assumption that $\tilde{E} = \emptyset$. This corresponds to the sizable assumption that all observations across all subjects share an identical conditional independence structure. Conversely, the approach of estimating a functional connectivity network for each subject independently corresponds to the assumption that $E^{pop} = \emptyset$ (or equivalently that $\Theta^{pop} = \boldsymbol{0}_{p \times p}$). In such a scenario, there is no advantage to be gained by leveraging information across subjects. Typically, we would expect the true underlying network structure across subjects to lie somewhere along the spectrum between these two extremes; thus justifying the proposed model.

3.2. *Mixed Neighborhood Selection.* In this section we formally describe the proposed MNS method. The proposed method can be seen as an extension of neighborhood selection to the *replicated* data scenario. The objective of such an extension is to adequately model the inter-subject heterogeneity present in fMRI data.

With this in mind we note that the linear regression approach described in Section 2.1.1 is particularly advantageous in our context where we have multiple related (i.e., clustered) longitudinal observations. Traditionally, such data has been modeled using linear mixed effects models [Pinheiro and Bates, 2000]. By adequately modeling the noise structure of clustered data, linear mixed models are able to effectively separate the population (i.e., fixed) effects from the subject-specific (i.e., random) effects. In doing so, they provide a highly flexible tool through which to analyze and describe many diverse datasets. They are therefore an ideal candidate in our scenario.

We consider learning the neighborhood of node $v \in V$ over a cohort of N subjects. We look to extend linear mixed effects models to the neighborhood selection scenario by proposing the following linear mixed effects model:

$$X_v^{(i)} = X_{\setminus v}^{(i)} \beta^v + X_{\setminus v}^{(i)} \tilde{b}^{(i),v} + \epsilon^{(i),v} \quad \text{for } i = 1, \dots N \ . \tag{3.2}$$

Recall that $X_v^{(i)}$ denotes the time series at node v for subject i. The model described in equation (3.2) directly extends traditional neighborhood selection model by introducing random effect terms, $\tilde{b}^{(i),v}$, for each subject. We note that β^v corresponds to the shared population neighborhood.

The random effects are assumed to follow a multivariate Gaussian distribution, $\tilde{b}^{(i),v} \sim \mathcal{N}(0, \Phi^v)$, independently of $\epsilon^{(i),v}$. Under such a model the the time series at each node thereby follows a Gaussian distribution of the following form:

$$X_v^{(i)} \sim \mathcal{N}\left(X_{\setminus v}^{(i)} \beta^v, X_{\setminus v}^{(i)} \Phi^v \left(X_{\setminus v}^{(i)} \right)^T + \sigma^2 I \right). \tag{3.3}$$

The choice of covariance structure for random effects is crucial to both estimating the model as well as to its interpretability. While it is possible to motivate many choices for $\Phi^v \in \mathbb{R}^{p-1 \times p-1}$, in this work we limit ourselves to the scenario where $\Phi^v = \sigma^2 \text{diag}(\sigma^{v2})$. Here $\sigma^v \in \mathbb{R}^{p-1}$ is a vector describing the standard deviation of the neighborhood of v across the cohort of N subjects. For example, a large σ_u^v value is indicative of a large variability in the edge between nodes v and u.

For any node $v \in V$, the model described in equations (3.2) and (3.3) is easily interpretable. The population, or fixed effects, neighborhood is captured in β^v. These are the effects that are shared across the entire cohort of subjects and correspond to the set of edges in E^{pop}. Meanwhile, the random effects are able to capture subject-specific deviations from the population neighborhood and can thereby be employed to obtain a network for each subject. Formally, the random effects captured in σ^v correspond to the set of highly variable edges, $\tilde{E}$. Finally, we are also able to obtain estimates of $\tilde{b}^{(i),v}$, which can be employed to obtain subject-specific networks. These values correspond to the subject-specific idiosyncrasies, $\tilde{E}^{(i)}$.

3.3. *Estimation algorithm.* The model described in equation (3.2) contains the following parameters, $\phi^v = (\beta^v, \sigma^v, \sigma^2) \in \mathbb{R}^{2(p-1)+1}$, which must be estimated. Given ϕ^v we can subsequently obtain the best linear unbiased predictions (BLUPs) for each of the random effects, $\tilde{b}^{(i),v}$, across subjects [Pinheiro and Bates, 2000]. In this work ϕ^v is estimated in a maximum likelihood framework. Following from equation (3.3), the negative log-likelihood for node v is proportional to:

$$\mathcal{L}(\phi^v) = \sum_{i=1}^{N} \frac{1}{2} \log \det V_v^{(i)} + \frac{1}{2} \left(X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v \right)^T V_v^{(i)^{-1}} \left(X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v \right), \tag{3.4}$$

where we define $V_v^{(i)}$ to be the variance structure for node v at subject i:

$$V_v^{(i)} = \sigma^2 \left(X_{\setminus v}^{(i)} \text{diag}(\sigma^{v2}) \left(X_{\setminus v}^{(i)} \right)^T + I \right). \tag{3.5}$$

In order to simplify future discussion, we re-parameterize the random effects component of the mixed effect model described in equation (3.2). We re-define the random effects term as follows:

$$\tilde{b}^{(i),v} = \text{diag}(\sigma^v)\, b^{(i),v}, \tag{3.6}$$

where $b^{(i),v} \sim \mathcal{N}(0, \sigma^2 I)$. Both the left and right sides of equation (3.6) follow the same distribution, however, this will simplify discussion of the optimization algorithm below.

While it would be possible to maximize equation (3.4) by obtaining the profile likelihood, in this work we treat the random effects as latent variables and employ an EM algorithm [McLachlan and Krishnan, 2007]. Fitting linear mixed effects models in this manner is a popular approach first posited by Dempster, Laird and Rubin [1977] and for which many efficient algorithms have been proposed [Meng and Van Dyk, 1998]. In the context of this work, such an approach will prove beneficial when regularization constraints are enforced on some elements of ϕ^v. Assuming the random effects, $b^{(i),v}$, are observed the complete data log-likelihood is proportional to:

$$\mathcal{L}_c(\phi^v) = \sum_{i=1}^{N} \frac{n+p}{2} \log \sigma^2 + \frac{1}{2\sigma^2} \left(\left|\left| X_v^{(i)} - X_{\setminus v}^{(i)} \beta^v - X_{\setminus v}^{(i)} \text{diag}(\sigma^v) b^{(i),v} \right|\right|_2^2 + b^{(i),v^T} b^{(i),v} \right). \tag{3.7}$$

It is important to note from equation (3.7) that the number of parameters remains fixed even as the number of subjects, N, increases. Therefore if a sufficiently large cohort of subjects is studied it is possible to estimate all parameters without the need to introduce regularization. However, regularization is introduced in this work for two reasons. First, sparse solutions remain feasible when only a reduced number of observations or subjects are available. Second, parsimonious solutions remain easily interpretable even in the presence of many nodes. As a result, we impose an ℓ_1 penalty on both the fixed as well as random effects. In terms of the random effects we penalize the variance terms, σ^v. If a variance is shrunk to zero, the resulting random effect is effectively removed from

the model. The introduction of sparsity inducing penalties yields the following penalized complete-data log-likelihood:

$$\mathcal{L}_c^{\lambda_1,\lambda_2}(\phi^v) = \mathcal{L}_c(\phi^v) + \lambda_1||\beta^v||_1 + \lambda_2||\sigma^v||_1, \tag{3.8}$$

where λ_1 and λ_2 are regularization parameters. Sparsity at the population level is enforced by λ_1, while λ_2 encourages sparsity in the random effects by shrinking the standard deviation terms, σ^v.

An EM algorithm is employed to minimize the penalized conditional log-likelihood. This involves iteratively computing the conditional expectation of latent variables, $Q(\phi;\phi^v)$, in our case the random effects, and minimizing the expected conditional log-likelihood with respect to parameters ϕ^v.

The expectation step (E-step) can be computed in closed form as follows:

$$b^{(i),v} = \left(\text{diag}(\sigma^v)X_{\setminus v}^{(i)^T}X_{\setminus v}^{(i)}\text{diag}(\sigma^v) + I\right)^{-1} X_{\setminus v}^{(i)^T}\text{diag}(\sigma^v)\left(X_v^{(i)} - X_{\setminus v}^{(i)}\beta^v\right) \tag{3.9}$$

independently for each subject $i = 1,\ldots,N$. It is clear from equation (3.9) that if σ_u^v is shrunk to zero then the uth entry of $b^{(i),v}$ will also be zero for all subjects.

In the minimization step (M-step) the latent variables, $b^{(i),v}$, are assumed to be observed. We therefore learn (β^v,σ^v) by solving the following convex problem:

$$(\beta^v,\sigma^v) = \underset{(\beta^v\in\mathbb{R}^p,\sigma^v\in\mathbb{R}_+^p)}{\text{argmin}} \left\{\left|\left|X_v^{(i)} - X_{\setminus v}^{(i)}\beta^v - X_{\setminus v}^{(i)}\text{diag}(b^{(i),v})\sigma^v\right|\right|_2^2 + \lambda_1||\beta^v||_1 + \lambda_2||\sigma^v||_1\right\}. \tag{3.10}$$

We note that equation (3.10) is Lasso problem with distinct regularization parameters applied to the fixed and random effects components respectively. A vast range of efficient algorithms can be employed to solve equation (3.10). In this work gradient descent algorithms [Friedman et al., 2007] were employed. The motivation behind this choice was that due to the iterative nature of the EM algorithm employed in this work, a lasso problem must be solved at each iteration. It follows that while solutions from one iteration to the next will typically not be identical they will be relatively similar. As a result, considerable computational gains can be obtained by using past solutions as warm-starts. Gradient descent algorithms are particularly well-suited for such tasks.

The proposed EM algorithm therefore iterates between equations (3.9) and (3.10). This is performed iteratively for each node $v \in V$. Once this is complete, networks can be obtained by applying the AND/OR rules described in equation (2.5). The MNS procedure is described in Algorithm 1.

3.4. *Parameter tuning.* The proposed method requires the tuning of two regularization parameters which govern the nature of the estimated population and subject-specific networks respectively. Large values of λ_1 will lead to sparse networks at the population level. Conversely, selecting large λ_2 will penalize the variance of the random effects leading to sparse subject-specific contributions to covariance structure.

Moreover, in the class of models considered in this work each covariate can contribute to the fixed as well as random effect structure. This can potentially lead to problems

Algorithm 1: Mixed Neighborhood Selection

Input: Data across N subjects, $\{X^{(i)}\}$, regularization parameters, λ_1, λ_2.

begin
  for $v \in \{1, \dots, V\}$ **do**
    Define initial estimates: $\beta^v = \boldsymbol{0}$, $\sigma^v = \boldsymbol{1}$, $\sigma = 1$ and $b^{(i),v} = \boldsymbol{0}$
    while *not converged* **do**
      Update (β^v, σ^v) by solving equation (3.10) `// M-step`
      Estimate latent variables using equation (3.9) for $i \in \{1, \dots, N\}$ `// E-step`
    Store β^v, σ^v and $\left\{b^{(i),v}\right\}_{i=1}^{N}$
  $E^{pop} = \{(u, v) : \beta_u^v \neq 0 \text{ and } \beta_v^u \neq 0\}$
  $\tilde{E} = \{(u, v) : \sigma_u^v \neq 0 \text{ and } \sigma_v^u \neq 0\}$
  $\tilde{E}^{(i)} = \{(u, v) : b_u^{(i),v} \neq 0 \text{ and } b_v^{(i),u} \neq 0\}$
return E^{pop}, $\tilde{E}$ *and* $\tilde{E}^{(i)}$ *for* $i = 1, \dots, N$

regarding the interpretability of estimated models. For example, over-penalizing the fixed effects may lead to over-estimation of the random effect variances as compensation [Schelldorfer, Bühlmann and Van de Geer, 2011]. The choice of regularization parameters is therefore a delicate issue which must be handled with care.

While information theoretic methods such as the Bayesian Information Criterion (BIC) may be employed for the purpose of tuning regularization parameters, in this work we employ cross-validation (CV) [Arlot, 2010]. As such, the data is divided into K folds. For each fold, the data from the remaining $K - 1$ folds is employed to fit the penalized linear mixed model described in Section 3.2. The resulting model is them used to predict the unseen data and the mean-square error is noted. This procedure is repeated and the results are averaged across all nodes. The pair of parameters which minimize the mean-square error are subsequently selected.

3.5. *Relationship to previous work.* The problem of simultaneously estimating multiple graphical models has recently received considerable attention. This problem is particularly relevant in the context of neuroscience where multi-subject studies are commonplace. Moreover, the recent collaborative trend towards sharing large neuroimaging datasets serves as an additional motivation [Poldrack and Gorgolewski, 2014].

From a methodological perspective several solutions have been proposed. Varoquaux et al. [2010] propose to leverage information by imposing a common sparse support for the precision matrices of all subjects. This is achieved via the use of a separate group Lasso penalty on the each edge. A similar penalty is also employed by Wee et al. [2014] in the context of performing classification between normal controls and MCI subjects. The Joint Graphical Lasso (JGL), proposed by Danaher, Wang and Witten [2014], generalized the above methods. Under their proposed framework, any convex penalty can be applied on edges across subjects and they propose both group as well as a fused lasso penalties. Other related methods include the work of Guo et al. [2011] and Pierson et al. [2015].

Our proposed method holds a number of advantages over these preceding methods.

First, it provides a principled manner for uncovering a population network that reflects edges shared across all subjects. While this could be achieved by combining observations across all subjects, such an approach carries with it the tenuous assumption that observations are interchangeable across subjects and will almost certainly be violated in practice. Approaches such as the JGL readily obtain estimates for each subject, however, there is no clear method for obtaining a population network. Second, the proposed method also reports edges which demonstrate high variability across a cohort. As far as we are aware, this is the first method to achieve this. Third, as we demonstrate in the simulation study below, the proposed method is capable of accurately recovering networks on both a subject and population level. We attribute this to the covariance model described in Section 3.1; by accurately learning which edges correspond to population (i.e., fixed effects) edges and which correspond to subject-specific idiosyncrasies the proposed method is able to leverage information across subjects in a discriminative manner. Lastly, the MNS algorithm is easily amenable to parallelization, making it particularly suitable for large datasets.

However, we must also note some of the disadvantages of our approach. The most significant of which is that we require an EM algorithm to estimate our model and therefore do not enjoy the guarantees provided by the convexity of the previous methods such as the JGL. In addition to this, since we take a neighborhood selection approach the resulting network estimates are not precision matrices and are therefore not guaranteed to be positive definite.

Finally, a related problem is that of quantifying statistically significant differences across multiple populations. Narayan, Allen and Tomson [2015] provide a principled manner for uncovering such differences. Much like the our work, they incorporate random effects into their model; however they do so in the context of modeling the presence or absence of an edge in multiple bootstrap estimates of the connectivity networks. In this way they are able to perform hypothesis tests on an edge-by-edge basis [Narayan and Allen, 2013, Narayan, Allen and Tomson, 2015]. Multiple other methods have also been developed for quantifying differences across populations. Ginestet et al. [2014] develop a central limit theorem for graph laplacians, allowing them to derive pivotal quantities and formally test for differences in pairs of networks. Zalesky, Fornito and Bullmore [2010] propose a network-based statistic to control the family-wise error rate when performing mass-univariate testing across all edges.

4. Simulation Study. In this section we evaluate the performance of the proposed method using simulated data that is representative of real functional imaging data. We assess the empirical performance of the MNS algorithm in three distinct settings which correspond to correctly reporting the edge structure of the population, subject-specific and highly variable network edges respectively. The first task corresponds to correctly recovering E^{pop} while the second requires learning subject-specific edge structure, $E^{(i)}$, defined in equation (2.1). Finally, the task of recovering variable edges is equivalent to learning the set of variable edges, $\tilde{E}$.

4.1. *Network simulation.* In order to perform such a study we require a method through which to simulate population and subject-specific networks. While numerous algorithms have been proposed to generate random individual networks (e.g., the Erdős and Rényi [1959] model and the preferential attachment model of Barabási and Albert [1999]), there has been limited work on algorithms to simulate multiple related networks. Notably, there is no documented method through which to generate networks from a cohort of related subjects that demonstrate the characteristics observed in real fMRI data; namely a shared core structure which is reproducible across all subjects together with significant inter-subject variability in the remaining edges [Bullmore, 2009, Ashby, 2011, Poldrack, Mumford and Nichols, 2011].

In order to address this issue we propose a new method of simulating networks. The proposed algorithm is motivated by an exploratory data analysis of resting state fMRI data from healthy subjects within the ABIDE consortium dataset. We briefly outline the proposed algorithm in this section with further details provided in Appendix B.

The proposed network simulation method proceeds as follows: first a population edge set, corresponding to E^{pop}, is randomly sampled according to the preferential attachment model of Barabási and Albert [1999]. These edges constitute the core, reproducible connectivity structure which will be present across all subjects. The precision matrix, Θ^{pop}, corresponding to E^{pop} is then obtained by randomly sampling edge weights as described in Danaher, Wang and Witten [2014].

A set of variable edges, $\tilde{E}$, is then selected randomly according to the Erdős and Rényi [1959] model, possibly overlapping with E^{pop}. For each subject, edges in $\tilde{E}$ are included in the subject-specific network with some fixed probability $\tau \in [0, 1]$. The edges weights associated with $\tilde{E}^{(i)}$ are randomly sampled for each subject and stored in $\Theta^{(i)}$. We note that the special case where $\tau = 0$ yields an identical network for all subjects. On the other hand, setting $\tau = 1$ still results in inter-subject variability as all edges within $\tilde{E}$ will have varying weights across subjects. This method is summarized in Algorithm 2.

Algorithm 2: Generate population and subject-specific random networks

Input: Number of nodes p, number of subjects N, size of random effects network $e_{ran} = |\tilde{E}|$, a random effects edge probability $\tau \in [0, 1]$ and connectivity strength $r \in \mathbb{R}_+$

begin
  Simulate E^{pop} according to Barabási and Albert [1999] model
  Build Θ^{pop} by uniformly sampling edge weights from the interval $[-r, -\frac{r}{2}] \cup [\frac{r}{2}, r]$
  Simulate $\tilde{E}$ according to Erdős and Rényi [1959] model with e_{ran} edges
  for $i \in \{1, \ldots, N\}$ **do**
    for *each edge* (j, k) **do**
      if $(j, k) \in \tilde{E}$ **then**
        $E^{(i)} \leftarrow E^{(i)} \cup (j, k)$ with probability τ
    Randomly select edge weights and signs for $\Theta^{(i)}$
return $E^{pop}, \tilde{E}$, $\{E^{(i)}\}$ *and* Θ^{pop}, $\{\Theta^{(i)}\}$

Algorithm 2 was employed to simulate synthetic data for a cohort of $N = 10$ subjects.

The number of nodes was fixed at $p = 50$, the size of the random effect network at $e_{ran} = 20$ and the random edge probability was $\tau = 1$. Datasets for each subject were simulated according to:

$$X^{(i)} \sim \mathcal{N}\left(0, \left(PD\left(\Theta^{pop} + \Theta^{(i)}\right)\right)^{-1}\right), \tag{4.1}$$

where $PD(\cdot)$ is a function applied in order to ensure the resulting matrix is positive definite (see Appendix B). Data was simulated with varying numbers of observations per subject, $n \in \{50, 100, 200\}$. We note that in all cases the number of observations per subject is far below the amount that would be sufficient to obtain a maximum likelihood estimate of the precision matrix. As a result, it is crucial to effectively leverage information across subjects.

Simulating networks as described in Algorithm 2 is only one of many possible methods which could be employed. In order to provide a thorough and fair comparison an additional simulation was also performed where networks were simulated as described in Danaher, Wang and Witten [2014]. This simulation was proposed with the objective of providing empirical evidence regarding how accurately subject-specific networks could be reported. It is therefore not well suited for examining how reliably the population or variance networks can be reported. The results for this simulation are provided in Appendix C.

4.2. *Alternative models.* Throughout this simulation the performance of the MNS algorithm was benchmarked against the current state of the art in each of the three settings described above. In the case of estimating the population network, the *graphical lasso* (Glasso) [Friedman, Hastie and Tibshirani, 2008] was employed. Such an approach has been used extensively in the neuroimaging community to learn functional connectivity networks across populations [Smith et al., 2011, Varoquaux and Craddock, 2013]. Moreover, an approach based on resampling and randomization was also employed. This approach, which we refer to as the *Stability* approach, is outlined in Appendix A. We note that while this approach is inspired by the recently proposed R^3 method of Narayan, Allen and Tomson [2015], the objective here is different. Formally, the R^3 method addresses the issue of comparing covariance structure across two populations while this simulation study is based on only a single population.

The problem of estimating subject-specific functional connectivity networks has received considerable attention. In this simulation study we compare the performance of the proposed method with the two penalized likelihood methods presented in Varoquaux et al. [2010] and Danaher, Wang and Witten [2014]. Both methods are related, enforcing a group or a fused lasso penalty across edges respectively, and we refer to each as the JGL-Group or the JGL-Fused algorithms respectively. The Glasso algorithm is also employed in this context.

As far as we are aware there are no alternative methods available which address the problem of recovering highly variable edges. In order to provide a benchmark for the MNS algorithm, the aforementioned *Stability* approach was employed in this context.

4.3. *Performance measures.* Throughout this simulation the task of recovering covariance structure is treated as a binary classification task. Thus performance is measured according to the proportion of edges which are correctly reported as being either present or absent. In order to compare performance across various algorithms we employ receiver operating characteristic (ROC) curves, which illustrate the performance of a binary classifier by plotting the true positive rate against false positive rate across a range of regularization parameters [Krzanowski and Hand, 2009].

The use of ROC curves requires a single, sparsity-inducing parameter to be varied across a range of possible values. As discussed previously, in the case of the MNS algorithm both the population and subject-specific parameters can affect sparsity. As a result, we look to reparameterize the MNS penalty as follows:

$$\lambda_1 = \alpha\lambda \tag{4.2}$$

$$\lambda_2 = \sqrt{2}(1-\alpha)\lambda \tag{4.3}$$

where α controls the ratio of sparsity between the population and subject-specific contributions and λ the overall sparsity. Thus α is fixed allowing λ to vary. While no such adjustments are needed in the case of the JGL-Fused algorithm, we follow the same parameterization described in equations (4.2) and (4.3) in the case of the JGL-Group algorithm[2].

4.4. *Simulation results.* In this section we present the results to the simulation study described above. We begin by first considering performance in the context of recovering the set of variable edges. This is a fundamental problem in neuroscientific applications [Kelly et al., 2012], however to date it has received limited attention. Results for the more frequently studied problems of recovering population and subject covariance structure are presented in Section 4.4.2 and 4.4.3 respectively.

Throughout this simulation the MNS algorithm was run with $\alpha = 0.25$ while sparsity parameter λ varied as described in equations (4.2) and (4.3). The same parameterization was employed for the JGL-Group algorithm with $\alpha = 0.15$ selected. In the case of the JGL-Fused algorithm, $\lambda_2 = 0.2$ was employed. Finally, the *Stability* algorithm was run with $B = 10,000$ bootstrap iterations per subject and $c = 0.25$.

4.4.1. *Variable network recovery.* Understanding variability in covariance structure across a cohort of subjects is a fundamental problem in neuroscience. In particular, understanding whether this variation can be attributed to phenotypic characteristics or other sources of noise (e.g., physiological noise or scanner noise) is crucial in further understanding the architecture of the human connectome.

The results shown in the top panel of Figure [2] demonstrate that the proposed MNS algorithm is able to accurately identify edges which demonstrate variability across a cohort of subjects. This is in contrast to the *Stability* method. Briefly, the *Stability* method (described in further detail in Appendix A) treats the presence or absence of

[2]Note this same parameterization was employed in the original study described in Danaher, Wang and Witten [2014]

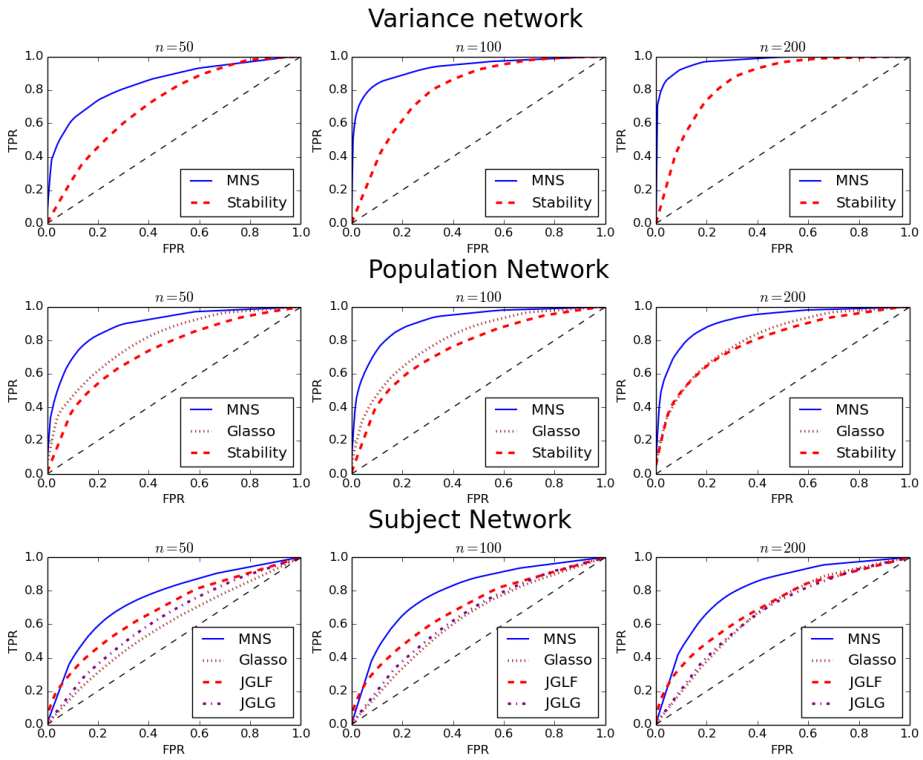

FIG 2. *Simulation results for all five algorithms across all tasks. Recovery of variable edges is considered in the top panel, population network recovery in shown in the middle panel and finally the bottom panel shows subject-specific network recovery. This simulation was performed with* $p = 50$ *nodes and* $n \in \{50, 100, 200\}$ *observations.*

edges at a subject level as a Bernoulli random variable. A hierarchical random effects model is then proposed to leverage information across all subjects. The resulting estimate of the edge variability, given in equation (A.6), is then employed to discriminate between variable and non-variable edges. The *Stability* method therefore corresponds to a two-step procedure where information is only leveraged *after* networks have been estimated for subjects independently. This is in contrast to the proposed method where subject-specific, population and variable networks are learnt *simultaneously*. This allows the MNS algorithm to effectively leverage information and results in significant improvements in performance. Further results are given in Table [1] where the true positive rate (TPR) and false positive rate (FPR) are reported for selected regularization parameters.

4.4.2. *Population network recovery.* Obtaining an accurate understanding of a population level covariance structure is a challenging problem due to the high inter-subject variability. As mentioned previously, it is imperative to differentiate between subject-specific idiosyncrasies and behavior which is reproducible across the entire cohort. A

popular approach often taken in neuroimaging studies is to estimate a single network using data from all subjects, thus effectively concatenating individuals' data. This corresponds to the sizable assumption that $\tilde{E} = \emptyset$. This approach is included in this simulation together with the aforementioned *Stability* approach. Here the mean estimate for edge presence, given by equation (A.5) was employed.

Results are shown in the middle panel of Figure [2]. It is interesting to note that for small sample sizes (i.e., $n = 50$ or $n = 100$) the *Stability* approach is out-performed by the Glasso. As mentioned in Section 4.4.1, we attribute this drop performance to its two-step design whereby information is only leveraged *after* networks have been estimated. It is only when the number of observations increases that reliable estimates of uncertainty can be obtained. Conversely, the difference in performance between the Glasso algorithm and the MNS algorithm is due to the presence of variable edges, implying $\tilde{E} \neq \emptyset$. Thus, by providing a more sophisticated model for inter-subject variability, the MNS algorithm is able to obtain more reliable population network estimates.

4.4.3. *Subject-specific network recovery.* Finally, we consider recovering subject-specific networks. This problem has received considerable attention in recent years and a range of methods have been proposed. The underlying theme in these methods revolves around effectively leveraging information across subjects. This is particularly evident in the JGL algorithm proposed by Danaher, Wang and Witten [2014] and the method presented by Varoquaux et al. [2010]. However, a short-coming of the aforementioned methods is that they leverage information in an indiscriminate manner. By enforcing either a group or fused lasso penalty on the entries of precision matrices, such methods effectively encourage homogeneous information leveraging across all edges.

While such approaches constitute significant methodological improvements, we envisage a scenario where edges can be ordered according to their variability (or reproducibility). This is a well-documented phenomenon in neuroimaging. In particular for fMRI data there is compelling evidence to suggest that variability in functional connectivity is directly modulated by the distance between regions [Power et al., 2012, Satterthwaite et al., 2012, Van Dijk, Sabuncu and Buckner, 2012].

The proposed MNS algorithm is able to address precisely this issue. By discriminating between subject-specific and population edges, it is able to effectively vary the how extensively each edge is leveraged across a population. As a result, the MNS algorithm is able to more reliably recover subject-specific covariance structure, as shown in the bottom panel of Figure [2].

5. Application. In this section the proposed MNS algorithm is applied to resting-state fMRI data from the ABIDE dataset [Di Martino et al., 2014]. While the ABIDE dataset contains data corresponding to healthy subjects and autism spectrum disorder (ASD) subjects, we chose only to study healthy controls here as the focus of this work consisted in fully understanding uncertainty across a single population of subjects. The decision to study the ABIDE dataset in this manner was motivated by the fact that it is an open-source dataset which has been widely studied, albeit often in the context of autism. While it would have been possible to study resting-state data for healthy

Algorithm	n	Population		Subject		Variance	
		TPR	FPR	TPR	FPR	TPR	FPR
MNS	50	0.76	0.12	0.75	0.33	0.54	0.06
	100	0.77	0.11	0.80	0.32	0.70	0.03
	200	0.75	0.11	0.82	0.30	0.79	0.02
Glasso	50	0.69	0.27	0.88	0.83	NA	
	100	0.70	0.27	0.83	0.66		
	200	0.68	0.27	0.85	0.58		
Stability	50	0.56	0.20	NA		0.54	0.24
	100	0.59	0.20			0.64	0.18
	200	0.78	0.35			0.71	0.15
JGL Group	50	NA		0.86	0.71	NA	
	100			0.83	0.62		
	200			0.82	0.57		
JGL Fused	50	NA		0.78	0.51	NA	
	100			0.79	0.51		
	200			0.79	0.50		

TABLE 1

Performance of all five algorithms considered. Results are presented for each of the three tasks: recovering population, subject and variance networks. For each algorithm the true positive rate (TPR) and false positive rate (FPR) are reported where applicable.

subjects across all sites, only the data from the University of Utah, School of Medicine (USM) site was considered here. This consisted of 58 subjects with ages ranging from 11 to 45 years old.

5.1. *Data acquisition and processing.* Resting-state fMRI data for healthy subjects was collected from the USM site [Di Martino et al., 2014], resulting in a cohort of 58 subjects. Whole-brain functional images were acquired by a 3T Siemens Magnetom TrioTim using an EPI sequence (T2*-weighted gradient echo; TR=2000 ms; TE=28 ms; flip angle=90°; 40 interleaved slices; voxel size= $3.4 \times 3.4 \times 3.0$ mm). In addition, structural images of each subject were acquired using a 3D-MPRAGE sequence (T1-weighted gradient echo TR= 2300 ms; TE = 2.91 ms; TI = 900 ms; flip angle = 9°; 160 sagittal slices; matrix size = 240×256; FOV = 256 mm; slice thickness = 1.2 mm). Preprocessing of the structural images involved removal of background noise and segmentation using using FAST [Zhang, Brady and Smith, 2001] and FIRST [Patenaude et al., 2011]. This resulted in 92 anatomical regions of interest.

Preprocessing of functional images involved motion correction using MCFLIRT and bandpass filtering (0.01-0.1 Hz) to extract low frequency fluctuations. In addition, a spatial single subject ICA was performed using FSL MELODIC [Smith et al., 2004]. Resulting independent components were automatically identified as either signal or noise using FSL FIX [Salimi-Khorshidi et al., 2014]. Functional data was subsequently denoised by regressing out noise independent components. Finally, motion parameters were regressed out. The first five volumes of each cleaned scan where subsequently discarded in order to allow for signal equilibration. Mean time-courses were extracted from all 92 regions, resulting in datasets with 230 observations over 92 nodes for each subject.

It follows that the number of observations available is far lower than the number re-

quired to accurately estimate the covariance structure for each subject; thereby implying that efficient leveraging of information across subjects is crucial.

5.2. *Results.* The MNS algorithm requires the specification of two regularization parameters, each of which controls the population and subject-specific topology of each node. As discussed in Section 3.4, the choice of these parameters is crucial to the interpretation of the resulting model. Parameter were selected on the basis of a 10-fold cross-validation framework.

One of the advantages of the proposed MNS algorithm is that it is able to simultaneously estimate both a population network, corresponding to reproducible edges which are present across the entire cohort of subjects, as well as a network quantifying variability on an edge-by-edge basis. The latter network is able to succinctly summarize variability across a cohort of subjects on an edge-by-edge basis.

On the top panel of Figure [3], the estimated population network is shown. The estimated network has an estimated edge density of around 10% and we note there is strong inter-hemispheric coupling as would be expected in resting-state connectivity. More importantly, the bottom panel of Figure [3] shows the estimated variability across all edges where the edge thickness is proportional to the estimated variance of the random effect. A hallmark of this plot is the anatomical distribution of highly variable edges. It has been suggested within the neuroimaging literature that head-movement may lead to the presence of spurious noise structures in fMRI data [Van Dijk, Sabuncu and Buckner, 2012, Satterthwaite et al., 2012, Power et al., 2012, Patel et al., 2014]. In particular, it has been suggested that subject motion in the scanner results in decreased estimated correlations between spatially remote regions and increased estimated correlations across spatially adjacent regions [Power et al., 2012]. Our results serve as a corroboration of this hypotheses as we discuss below.

Figure [4] plots the estimated variance against euclidean distance for each edge. We note there is a quadratic effect, with the large variances exhibited between spatially proximal and remote regions. This is precisely the effect reported in various neuroimaging studies (e.g., Satterthwaite et al. [2012] and Power et al. [2012]). Following from the aforementioned works, we hypothesize that this variability may have been introduced by the individual head motion of each subject.

It is also interesting to study the localization of highly variable edges in the bottom panel of Figure [3]. In particular, we note there is high variability within both the prefrontal and orbitofrontal regions; this is clearly seen in the sagittal and axial plots. These regions are known to be particularly problematic [Deichmann et al., 2003] due to their proximity to air/tissue interfaces which leads to image distortion and signal losses and have been reported as being highly variable across subjects [Finn et al., 2015, Miranda-Dominguez et al., 2014]. It is therefore reassuring that the proposed method is able to report such variability in order to obtain more reliable networks at the population and subject-specific level.

6. Discussion. Understanding variability in functional connectivity across a population of subjects remains a fundamental challenge with important consequences in

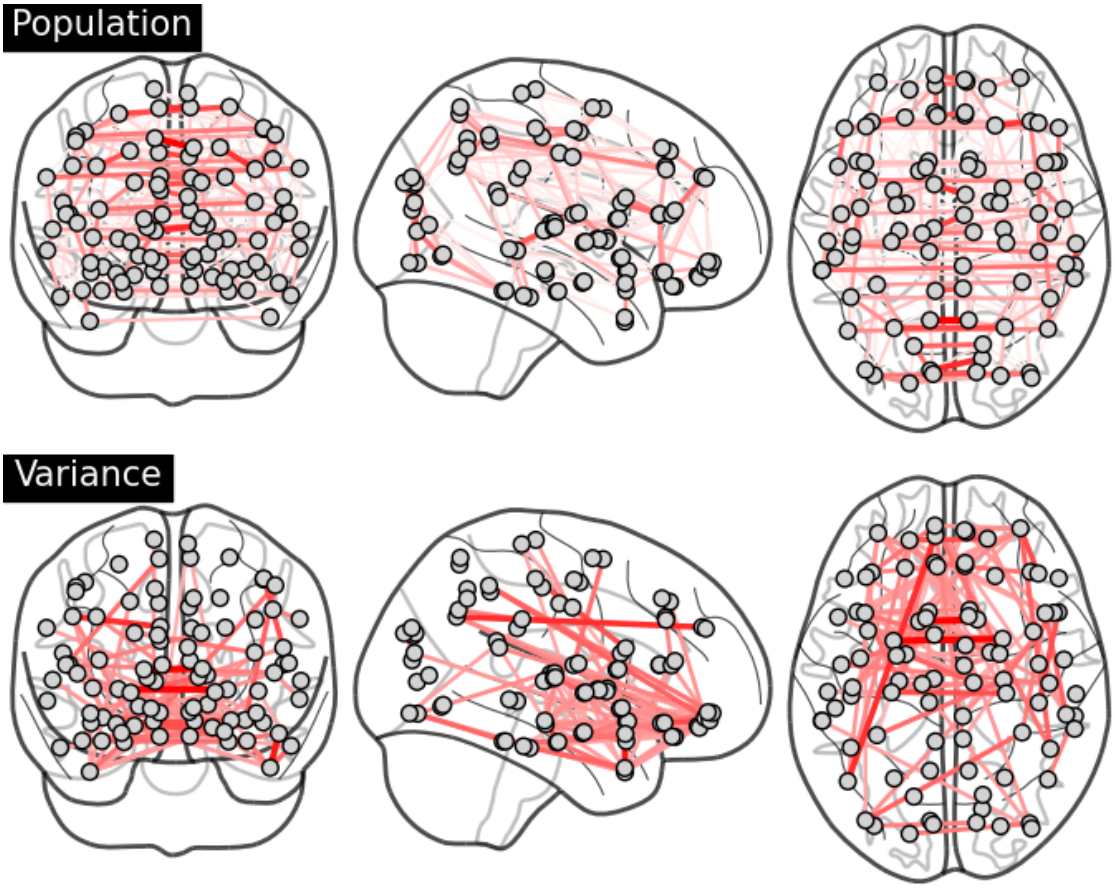

FIG 3. *Estimated population network (top) and variable edge network (bottom). Edges color indicates positive (red) or negative (blue) partial correlations and thickness is proportional to the magnitude of the partial correlation (or variance in the case of variance network). In the case of the population network we note there is high inter-hemispheric coupling which is to be expected in resting-state data. Conversely, the variable edge network demonstrates interesting spatial structure; highly variable edges occur between spatially proximal or remote regions as often hypothesized within the neuroimaging community. We also note that variability seems to conglomerate near the prefrontal and orbitofrontal regions.*

modern neuroimaging. Formally, understanding variability is seen as a crucial stepping stone towards obtaining a more holistic understanding and interpretation of functional connectivity [Kelly et al., 2012, Fallani et al., 2014]. This in turn will naturally lead to a better understanding of many diseases which are thought to be caused or characterized by changes in functional connectivity [Uddin, Supekar and Menon, 2013]. From a statistical perspective estimating variability in second order statistics remains challenging due to positive-definiteness and identifiability constraints. Mixed effects models are ideally suited in the context of regression [Pinheiro and Bates, 2000], and have been extend to the domain of GGMs here by leveraging the results of Meinshausen and Bühlmann [2006].

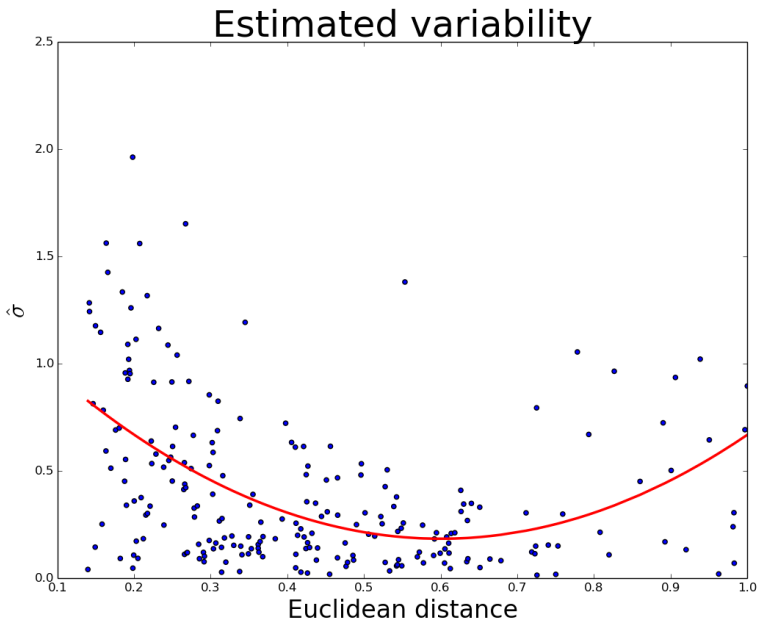

FIG 4. *Estimated edge variability as a function of (scaled) Euclidean distance between regions. The quadratic line of best fit in plotted in red. This serves to highlight our claim that variability is highest between spatially adjacent or remote regions.*

MNS looks to accurately learn the set of highly variable edges across a cohort of subjects whilst simultaneously obtaining more reliable network estimates both at the population and subject-specific level. This is achieved by proposing a more sophisticated model for the underlying covariance structure. The proposed model looks to decompose covariance structure as the union of population effects which are reproducible across subjects with subject-specific idiosyncrasies. In this way, the MNS algorithm is able to leverage information across subjects in a discriminative manner. This is in contrast to many of the current methodologies which leverage information in an indiscriminate fashion (e.g., via the use of group or fused regularization penalties whose parameterization is fixed across edges).

Throughout a series of simulation studies we demonstrate that MNS is able to accurately recover population and subject-specific functional connectivity networks. The MNS algorithm is also shown to be capable to recovering the set of highly variable edges. The simulations also provide empirical evidence of our claim that improved estimates of connectivity networks can be obtained by leveraging information across subjects in a discriminate manner. Moreover, a novel algorithm for simulation multiple related networks is proposed for the simulations. While there is a wide number of algorithms available to simulate individual networks [Rubinov and Sporns, 2010], the question of simulating multiple related networks has received limited attention.

The MNS algorithm requires the specification of two regularization parameters, λ_1 and

λ_2, each of which has a natural interpretation. The first parameter controls the sparsity in the population node topologies while the second controls the sparsity of the subject-specific edges. We propose to learn both parameters via cross-validation. Although the use of cross-validation incurs an additional computational cost, our experience is that it results in more accurate estimates of functional connectivity through improved selection of regularization parameters. However, a shortcoming of the proposed method is that both regularization parameters are dependent on each other [Schelldorfer, Bühlmann and Van de Geer, 2011]. Thus a large choice of λ_1 can be offset by a small choice of λ_2, leading fixed effects to be explained by inflated random effects estimates. An additional shortcoming is that estimated functional connectivity networks must be inferred via the use of an AND or OR rule [Meinshausen and Bühlmann, 2006]. While such an approach is able to recover the sparse support structure, it is not able to obtain an estimate of the corresponding sparse precision matrix which alternative methods such as the JGL are capable of obtaining. Despite these shortcomings, we believe the additional richness of information provided by the MNS algorithm will make it invaluable in further understanding functional connectivity.

The MNS algorithm, together with network simulation methods described in this work have been implemented as an R package named `MNS`. This can be downloaded from the Comprehensive R Archive Network (CRAN).

MNS was applied to resting-state fMRI data taken from the ABIDE consortium [Di Martino et al., 2014]. The results serve to corroborate a widely held hypothesis within the neuroimaging community regarding the effect of subject movement of functional connectivity [Power et al., 2012, Satterthwaite et al., 2012] and are consistent with the literature [Deichmann et al., 2003]. As far as we are aware, this is the first work to formally quantify variability in functional connectivity across a cohort of subjects. In future, MNS could also be employed to study variability across populations including healthy controls and subjects suffering from certain neuropathologies. Furthermore, another interesting avenue would be to employ MNS to study variability during task-based experiments. It has been suggested that connectivity in this context is non-stationary [Hutchison et al., 2013, Monti et al., 2014, 2015]. Extending MNS to this scenario would allow for the study of variability induced by cognitive tasks.

In conclusion, the MNS algorithm provides a method for simultaneously estimating population and subject-specific functional connectivity whilst also reporting highly variable edges across a population.

APPENDIX A: STABILITY APPROACH FOR *REPLICATED* DATA

In this section we briefly overview a stability selection (i.e., bootstrap) approach for studying *replicated* data. This approach is inspired by the R^3 approach proposed in Narayan, Allen and Tomson [2015], however, it has a fundamentally different objective. As a result, some adjustments are introduced.

As in the R^3 method, this approach is based upon resampling, randomized penalizations and random effects. The method, described in Algorithm 3, proceeds by iteratively obtaining bootstrapped estimates of covariance structure for each subject. These results are subsequently incorporated into a Beta-Binomial random effects model. Each of these steps is described below, for further discussion and motivation of these steps we refer the reader to Narayan and Allen [2013] and Narayan, Allen and Tomson [2015].

A.1. Resampling. In order to obtain reliable estimates of covariance structure the bootstrap is employed; resulting in B bootstrap estimates of connectivity structure per subject. Recall that the dataset for the ith subject, $X^{(i)} \in \mathbb{R}^{n\times p}$, consists of n observations across p nodes. At the bth bootstrap iteration, n observations are sampled with replacement in order to form a bootstrapped dataset, $X^{(i),b} \in \mathbb{R}^{n\times p}$, which is subsequently used to obtain an estimate for the covariance structure using the Graphical Lasso, as described in Section A.2.

A.2. Randomization penalization. In order to alleviate possible bias introduced by the use of an ℓ_1 penalty we employ randomized penalization techniques, an approach first introduced by Meinshausen and Bühlmann [2010]. The objective of randomized penalization schemes is to reduce the influence of inclusion/exclusion of any edge on the presence of remaining edges. Thus when estimating the network for the bth bootstrap sample, a random, symmetric penalty matrix, $\Lambda^{(i),b} \in \mathbb{R}^{p\times p}$, is employed.

In order to obtain $\Lambda^{(i),b}$, we first estimate the regularization parameter for the ith subject using the StARS method of Liu, Roeder and Wasserman [2010]. This is performed only once using the entire (non-bootstrapped) dataset, $X^{(i)}$, and is denoted by $\lambda^{(i)} \in \mathbb{R}$. The randomized penalization matrix is defined as follows:

$$\left(\Lambda^{(i),b}\right)_{k,j} = \left(\Lambda^{(i),b}\right)_{j,k} = \lambda^{(i)} + c\lambda^{(i)}_{max} W_{j,k} \quad \forall j < k, \tag{A.1}$$

where $\lambda^{(i)}_{max}$ is the value of sparsity parameter leading to a null model and $W \in \{-1,+1\}^{p\times p}$ is defined as:

$$W_{j,k} = \begin{cases} +1, & \text{w.p. } 0.5 \\ -1, & \text{w.p. } 0.5 \end{cases}.$$

We are then able to obtain a penalized estimate of the precision as follows:

$$\Theta^{(i),b} = \underset{\Theta}{\operatorname{argmin}} \left\{ -\log\det\Theta + \operatorname{trace}\left(\frac{1}{n} X^{(i),b^T} X^{(i),b} \Theta\right) + ||\Lambda^{(i),b} \circ \Theta||_1 \right\}, \tag{A.2}$$

where $\circ$ denotes element-wise multiplication.

A.3. Random effects. Finally, we look to formally quantify the presence or absence of edges at a population level. In order to achieve this a Beta-Binomial model is employed. For the ith subject we treat the presence of any given edge at each bootstrap iteration as a Binomial random variable. We thus define $Y^{(i),B} \in \mathbb{R}^{p\times p}$ such that

$$Y^{(i),B}_{j,k} = \frac{1}{B}\sum_{b=1}^{B} \mathbb{I}\left(\Theta^{(i),b}_{j,k} \neq 0\right). \tag{A.3}$$

Following Narayan, Allen and Tomson [2015], $Y^{(i),B}_{j,k}$ is modeled as follows:

$$Y^{(i),B}_{j,k} | \mu^{(i)}_{j,k} \sim \text{Binomial}\ (\mu^{(i)}_{j,k}, B) \quad \text{and} \quad \mu^{(i)}_{j,k} \sim \text{Beta}\ (\mu^{pop}_{j,k}, \rho^{pop}_{j,k}), \tag{A.4}$$

where $\mu^{pop}_{j,k}$ is the population mean and $\rho^{pop}_{j,k}$ the variance. They can be estimated as follows:

$$\mu^{pop}_{j,k} = \frac{1}{N}\sum_{i=1}^{N} Y^{(i),B}_{j,k} \tag{A.5}$$

$$\rho^{pop}_{j,k} = \frac{B}{B-1}\,\frac{\sum_{i=1}^{N}\left(\mu^{pop}_{j,k} - Y^{(i),B}_{j,k}\right)^2}{\mu^{pop}_{j,k}(1-\mu^{pop}_{j,k})(N-1)} - \frac{1}{B-1} \tag{A.6}$$

These parameters are subsequently used to infer population networks (using μ^{pop}) as well as report highly variable edges (using ρ^{pop}). Pseudo-code for the *Stability* approach is provided in Algorithm 3.

Algorithm 3: Stability approach for *replicated* data

Input: Data across N subjects, $\{X^{(i)}\}$, number of bootstrap samples to perform, B.

```
begin
    for i ∈ {1, ..., N} do
        Select λ^(i) using the StARS method [Liu, Roeder and Wasserman, 2010]
        for b ∈ {1, ..., B} do
            Obtain X^(i),b be sampling n times with replacement from X^(i)
            Set randomization penalization matrix, Λ^(i),b, as in equation (A.1)
            Estimate penalization precision matrix, Θ^(i),b, as in equation (A.2)
    Estimate μ^pop, ρ^pop using equations (A.5) and (A.6)
return μ^pop, ρ^pop
```

APPENDIX B: SIMULATING A COHORT OF NETWORKS

Producing synthetic data where the true underlying covariance structure is known is fundamental to providing an empirical validation of any algorithm. Moreover, in our scenario is it is crucial to ensure the simulated data demonstrates many of the properties often encountered in neuroimaging datasets. On a subject-specific level, this corresponds

to a power law distribution of edges across nodes together with the presence of hub nodes. Such networks can be efficiently simulated using the scale-free network algorithm of Barabási and Albert [1999].

However, the problem of simulating multiple networks for multiple related subjects has not been thoroughly considered in the literature. While there is a vast literature on the properties which can be expected for subject-specific networks (see e.g., Bullmore [2009]), there is limited knowledge of the behavior which can be expected across a cohort of related subjects. In this work we look to address this issue by empirically studying resting state data from healthy subjects taken from the ABIDE consortium [Di Martino et al., 2014].

The dataset employed here consisted of a resting state scan for $N = 58$ subjects. For each subject, $n = 230$ BOLD measurements were collected over $p = 92$ ROIs. We consequently estimated functional connectivity networks for each subject independently while employing a Graphical Lasso (Glasso) penalty. A stability selection procedure was employed, whereby a block bootstrap was used to resample the data for each subject multiple times. Randomized penalization was also employed to further correct for systematic bias in model selection. At each iteration, selected edges were recorded. This allowed us to obtain a network per subject by selecting only edges that where consistently present across all iterations. A population network was also obtained by selecting the edges that where consistently present across all subjects.

The properties of the resulting networks where studied in the hope of obtaining notable properties which could subsequently be recreated synthetically. In particular, we chose to focus on graph theoretic measures as these can be easily interpreted and calculated on simulated data [Rubinov and Sporns, 2010]. Specifically, the clustering coefficient was studied across all networks. This provides a measure of how tightly nodes in a network tend to group together and expresses the cohesiveness of a network [Barrat et al., 2004].

The results, shown in Figure [5], show that the clustering coefficient is significantly larger in the population network when compared the each of the subject-specific networks. We hypothesize that this is a manifestation of the fact that there is a highly structured population network present. We further hypothesize that it is the large inter-subject variability which accounts for some of the drop in clustering coefficient at the subject-specific level.

When simulating synthetic networks, as described in Algorithm 2, we look to recreate these properties (i.e., the drop in clustering coefficient). This is achieved by first simulating a population network according to the Barabási and Albert [1999] model. This results in a highly structure population network which also demonstrates many of the properties known to be present in neuroimaging data (e.g., power law distribution and the presence of hub nodes). A subset of highly variable edges, denoted by $\tilde{E}$, is then randomly selected according to the Erdős and Rényi [1959] model. For each subject, each edge in $\tilde{E}$ is added to the subject-specific network with a given probability, τ. This yields variable edges that are only present across a subset of the population. The introduction of these random edges serves to reduce the clustering coefficient of the subject-specific network, thereby recreating the properties observed in our exploratory analysis.

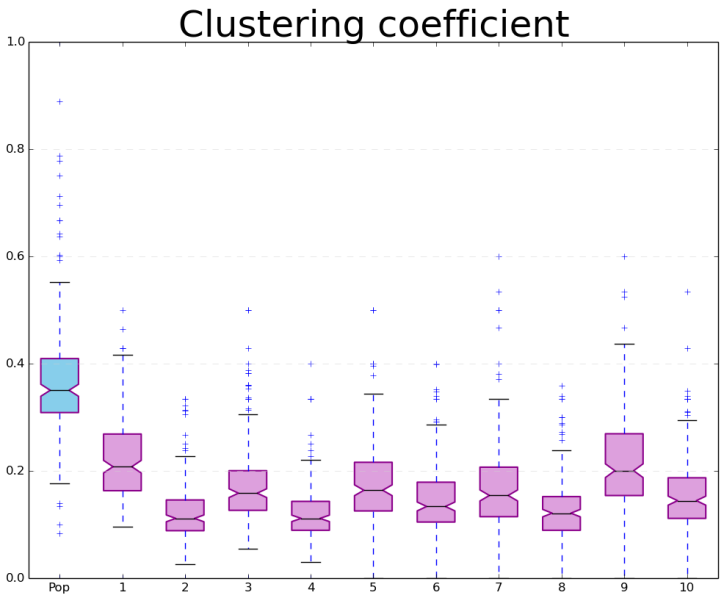

FIG 5. *Clustering coefficients for the population network (in blue) as well as for 10 randomly selected subjects. We note there is a clear drop in clustering coefficient from the population network to the subject-specific networks.*

B.1. Ensuring positive definiteness. Through algorithm 2 we are able to simulate a population precision, Θ^{pop}, together with subject-specific deviations, $\Theta^{(i)}$. We define the precision for each subject to be $\Theta^{pop} + \Theta^{(i)}$, however, care must be taken to ensure this sum is positive-definite. In this work we follow Danaher, Wang and Witten [2014] and ensure the subject-specific precision matrices are positive definite by rescaling the matrix. Formally, each off-diagonal element is divided by the sum of the absolute values of all off-diagonal elements in its row. This yields a non-symmetric matrix which is subsequently averaged with its transpose.

APPENDIX C: FURTHER SIMULATIONS

In Section 4 networks were simulated as described in Algorithm 2. While this algorithm was derived from an exploratory analysis of resting-state fMRI data, a wide range of alternative algorithms could also be proposed. In this section we look to provide further empirical evidence by recreating the simulation study of Danaher, Wang and Witten [2014].

While Danaher, Wang and Witten [2014] are able to simulate networks where variability is present, their proposed method is designed primarily to provide empirical evidence on how accurately subject-specific networks could be recovered. We therefore follow Danaher, Wang and Witten [2014] and focus exclusively on recovering subject-specific covariance structure here.

Two simulations where performed where data was simulated for $N = 3$ subjects and $p = 100$ and $p = 250$ nodes respectively. Within each simulation, nodes where divided into 10 equally sized and unconnected components. The connectivity structure within each component was simulated according to scale-free model of Barabási and Albert [1999], resulting in 10 scale-free sub-networks. Of the 10 sub-networks, eight were shared across all three subjects. Of the remaining two sub-networks, one was present in two out of three subjects and the final sub-network was only present in the first subject. For further see Danaher, Wang and Witten [2014].

C.1. Simulation 1: $p = 100$, $n \in \{50, 100, 200\}$. In the first simulation $p = 100$ nodes were employed resulting in 10 sub-networks each with 10 nodes. The number of observations per subject was allowed to vary from $n = 50$ through to $n = 200$. The results over 100 simulations are shown in top row of Figure [6]. The MNS algorithm performs competitively with respect to the JGL-Fused algorithm and outperforms both the JGL-Group and graphical lasso algorithms across all values of n. In particular, we note that the MNS algorithm remains competitive even as the number of observations, n, falls drastically.

C.2. Simulation 2: $p = 250$, $n \in \{50, 100, 200\}$. The second simulation employed $p = 250$ nodes which were divided into 10 sub-networks of 25 nodes each. The number of observations per subject was allowed to vary from $n = 50$ through to $n = 200$ as before. The results over 100 simulations are shown in bottom row of Figure [6]. As before, the MNS algorithm performs competitively against alternative algorithms. As with the previous simulation, we note there is a trend for ROC curves to improve as the number of observations, n, increases.

p	n	MNS		Glasso		JGL Fused		JGL Group	
100	50	0.346	0.006	0.175	0.016	0.343	0.007	0.221	0.012
	100	0.477	0.003	0.282	0.008	0.503	0.005	0.353	0.005
	200	0.594	0.002	0.429	0.004	0.632	0.005	0.514	0.003
250	50	0.287	0.002	0.125	0.007	0.292	0.003	0.161	0.005
	100	0.443	0.002	0.215	0.003	0.451	0.002	0.295	0.003
	200	0.573	0.001	0.370	0.002	0.584	0.002	0.465	0.002

TABLE 2

Performance of all four algorithms when recovering subject specific functional connectivity structure

ACKNOWLEDGEMENTS

The authors wish to acknowledge the help of Dr. Gareth Ball in preprocessing the data.

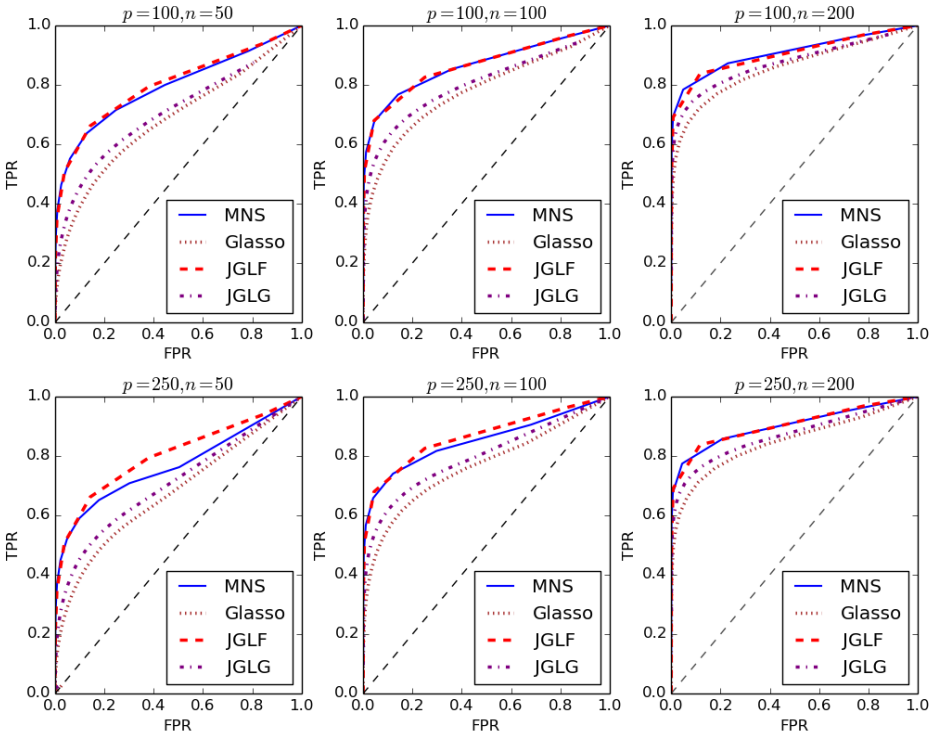

FIG 6. *Results for simulations 1 and 2. These simulations sought to re-create the simulation study presented in Danaher, Wang and Witten [2014] with* $p = 100$ *and* $p = 250$ *nodes respectively.*

REFERENCES

ARLOT, S. (2010). A survey of cross-validation procedures for model selection. *Statistics surveys* **4** 40–79.

ASHBY, G. (2011). *Statistical analysis of fMRI data*. MIT press.

BARABÁSI, A.-L. and ALBERT, R. (1999). Emergence of scaling in random networks. *science* **286** 509–512.

BARRAT, A., BARTHELEMY, M., PASTOR-SATORRAS, R. and VESPIGNANI, A. (2004). The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America* **101** 3747–3752.

BULLMORE, E. and SPORNS, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* **10** 186–198.

DAMOISEAUX, J., ROMBOUTS, S., BARKHOF, F., SCHELTENS, P., STAM, C., SMITH, S. and BECKMANN, C. (2006). Consistent resting-state networks across healthy subjects. *Proceedings of the national academy of sciences* **103** 13848–13853.

DANAHER, P., WANG, P. and WITTEN, D. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76** 373–397.

DEICHMANN, R., GOTTFRIED, J., HUTTON, C. and TURNER, R. (2003). Optimized EPI for fMRI studies of the orbitofrontal cortex. *Neuroimage* **19** 430–441.

DEMPSTER, A. (1972). Covariance selection. *Biometrics* 157–175.

Dempster, A., Laird, N. and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)* 1–38.

Di Martino, A., Yan, C., Li, Q., Denio, E., Castellanos, F., Alaerts, K., Anderson, J., Assaf, M., Bookheimer, S. and Dapretto, M. (2014). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry* **19** 659–667.

Erdős, P. and Rényi, A. (1959). On random graphs. *Publicationes Mathematicae Debrecen* **6** 290–297.

Fallani, F., Richiardi, J., Chavez, M. and Achard, S. (2014). Graph analysis of functional brain networks: practical issues in translational neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369** 20130521.

Finn, E., Shen, X., Scheinost, D., Rosenberg, M., Huang, J., Chun, M., Papademetris, X. and Constable, T. (2015). Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience*.

Friedman, J., Hastie, T. and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.

Friedman, J., Hastie, T., Höfling, H. and Tibshirani, R. (2007). Pathwise coordinate optimization. *The Annals of Applied Statistics* **1** 302–332.

Friston, K. (1994). Functional and effective connectivity in neuroimaging: a synthesis. *Human brain mapping* **2** 56–78.

Friston, K. (2011). Functional and effective connectivity: a review. *Brain connectivity* **1** 13–36.

Ginestet, C., Balanchandran, P., Rosenberg, S. and Kolaczyk, E. (2014). Hypothesis Testing For Network Data in Functional Neuroimaging. *arXiv preprint arXiv:1407.5525*.

Guo, J., Levina, E., Michailidis, G. and Zhu, J. (2011). Joint estimation of multiple graphical models. *Biometrika*.

Huettel, S., Song, A. and McCarthy, G. (2004). *Functional magnetic resonance imaging* **1**. Sinauer Associates Sunderland.

Hutchison, M., Womelsdorf, T., Allen, E., Bandettini, P., Calhoun, V., Corbetta, M., Della Penna, S., Duyn, J., Glover, G. and Gonzalez-Castillo, J. (2013). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* **80** 360–378.

Kanai, R. and Rees, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience* **12** 231–242.

Kelly, C., Biswal, B., Craddock, C., Castellanos, X. and Milham, M. (2012). Characterizing variation in the functional connectome: promise and pitfalls. *Trends in cognitive sciences* **16** 181–188.

Krzanowski, W. and Hand, D. (2009). *ROC curves for continuous data*. CRC Press.

Lauritzen, S. (1996). *Graphical models*. Oxford University Press.

Lee, H., Lee, D., Kang, H., Kim, B. and Chung, M. (2011). Sparse brain network recovery under compressed sensing. *Medical Imaging, IEEE Transactions on* **30** 1154–1165.

Liu, H., Roeder, K. and Wasserman, L. (2010). Stability approach to regularization selection (stars) for high dimensional graphical models. In *Advances in neural information processing systems* 1432–1440.

Markov, N., Ercsey-Ravasz, M., Van Essen, D., Knoblauch, K., Toroczkai, Z. and Kennedy, H. (2013). Cortical high-density counterstream architectures. *Science* **342** 1238406.

McLachlan, G. and Krishnan, T. (2007). *The EM algorithm and extensions* **382**. John Wiley & Sons.

Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* 1436–1462.

Meinshausen, N. and Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72** 417–473.

Meng, X. and Van Dyk, D. (1998). Fast EM-type implementations for mixed effects models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60** 559–578.

Miranda-Dominguez, O., Mills, B., Carpenter, S., Grant, K., Kroenke, C., Nigg, J. and Fair, D. (2014). Connectotyping: Model Based Fingerprinting of the Functional Connectome. *PLoS ONE*.

Monti, R., Hellyer, P., Sharp, D., Leech, R., Anagnostopoulos, C. and Montana, G. (2014).

Estimating time-varying brain connectivity networks from functional MRI time series. *Neuroimage* **103** 427–443.

Monti, R., Lorenz, R., Hellyer, P., Leech, R., Anagnostopoulos, C. and Montana, G. (2015). Graph embeddings of dynamic functional connectivity reveal discriminative patterns of task engagement in HCP data. In *Pattern Recognition in NeuroImaging (PRNI), 2015 International Workshop on* 1–4. IEEE.

Mueller, S., Wang, D., Fox, M., Yeo, T., Sepulcre, J., Sabuncu, M., Shafee, R., Lu, J. and Liu, H. (2013). Individual variability in functional connectivity architecture of the human brain. *Neuron* **77** 586–595.

Narayan, M. and Allen, G. (2013). Randomized Approach to Differential Inference in Multi-subject Functional Connectivity. In *Pattern Recognition in Neuroimaging (PRNI), 2013 International Workshop on* 78–81. IEEE.

Narayan, M., Allen, G. and Tomson, S. (2015). Two Sample Inference for Populations of Graphical Models with Applications to Functional Connectivity. *arXiv preprint arXiv:1502.03853*.

Patel, A., Kundu, P., Rubinov, M., Jones, S., Vértes, P., Ersche, K., Suckling, J. and Bullmore, E. (2014). A wavelet method for modeling and despiking motion artifacts from resting-state fMRI time series. *NeuroImage* **95** 287–304.

Patenaude, B., Smith, S., Kennedy, D. and Jenkinson, M. (2011). A Bayesian model of shape and appearance for subcortical brain segmentation. *Neuroimage* **56** 907–922.

Pierson, E., Koller, D., Battle, A., Mostafavi, S. and Consortium, G. (2015). Sharing and specificity of co-expression networks across 35 human tissues.

Pinheiro, J. and Bates, D. (2000). *Mixed-effects models in S and S-PLUS*. Springer Science & Business Media.

Pircalabelu, E., Claeskens, G., Jahfari, S. and Waldorp, L. (2015). Focused Information Criterion for Graphical Models in fMRI connectivity with high-dimensonal data. *Annals of Applied Statistics*.

Poldrack, R. and Gorgolewski, K. (2014). Making big data open: data sharing in neuroimaging. *Nature neuroscience* **17** 1510–1517.

Poldrack, R., Mumford, J. and Nichols, T. (2011). *Handbook of functional MRI data analysis*. Cambridge University Press.

Power, J., Barnes, K., Snyder, A., Schlaggar, B. and Petersen, S. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* **59** 2142–2154.

Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* **52** 1059–1069.

Ryali, S., Chen, T., Supekar, K. and Menon, V. (2012). Estimation of functional connectivity in fMRI data using stability selection-based sparse partial correlation with elastic net penalty. *Neuroimage* **59** 3852–3861.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C., Glasser, M., Griffanti, L. and Smith, S. (2014). Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* **90** 449–468.

Satterthwaite, T., Wolf, D., Loughead, J., Ruparel, K., Elliott, M., Hakonarson, H., Gur, R. and Gur, R. (2012). Impact of in-scanner head motion on multiple measures of functional connectivity: relevance for studies of neurodevelopment in youth. *Neuroimage* **60** 623–632.

Schelldorfer, J., Bühlmann, p. and Van de Geer, S. (2011). Estimation for High-Dimensional Linear Mixed-Effects Models Using l1-Penalization. *Scandinavian Journal of Statistics* **38** 197–214.

Smith, S. (2012). The future of FMRI connectivity. *Neuroimage* **62** 1257–1266.

Smith, S., Jenkinson, M., Woolrich, M., Beckmann, C., Behrens, T., Johansen-Berg, H., Bannister, P., De Luca, M., Drobnjak, I. and Flitney, D. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* **23** S208–S219.

Smith, S., Miller, K., Salimi-Khorshidi, G., Webster, M., Beckmann, C., Nichols, T., Ramsey, J. and Woolrich, M. (2011). Network modelling methods for FMRI. *Neuroimage* **54** 875–891.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 267–288.

Uddin, L., Supekar, K. and Menon, V. (2013). Reconceptualizing functional brain connectivity in autism from a developmental perspective. *Frontiers in human neuroscience* **7**.

Van Dijk, K., Sabuncu, M. and Buckner, R. (2012). The influence of head motion on intrinsic functional connectivity MRI. *Neuroimage* **59** 431–438.

Varoquaux, G. and Craddock, C. (2013). Learning and comparing functional connectomes across subjects. *NeuroImage* **80** 405–415.

Varoquaux, G., Gramfort, A., Poline, J. and Thirion, B. (2010). Brain covariance selection: better individual functional connectivity models using population prior. In *Advances in Neural Information Processing Systems* 2334–2342.

Wee, C., Yap, P., Zhang, D., Wang, L. and Shen, D. (2014). Group-constrained sparse fMRI connectivity modeling for mild cognitive impairment identification. *Brain Structure and Function* **219** 641–656.

Wehbe, L., Ramdas, A., Steorts, R. and Shalizi, C. (2015). Regularized Brain Reading with Shrinkage and Smoothing. *Annals of Applied Statistics*.

Zalesky, A., Fornito, A. and Bullmore, E. (2010). Network-based statistic: identifying differences in brain networks. *Neuroimage* **53** 1197–1207.

Zhang, Y., Brady, M. and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *Medical Imaging, IEEE Transactions on* **20** 45–57.

Zilles, K. and Amunts, K. (2013). Individual variability is not noise. *Trends in cognitive sciences* **17** 153–155.

Zuo, X., Di Martino, A., Kelly, C., Shehzad, Z., Gee, D., Klein, D., Castellanos, X., Biswal, B. and Milham, M. (2010). The oscillating brain: complex and reliable. *Neuroimage* **49** 1432–1445.

R. P. Monti, C. Anagnostopoulos
Department of Mathematics,
Imperial College London
SW7 2AZ, London, UK
E-mail: ricardo.monti08@imperial.ac.uk
canagnos@imperial.ac.uk

G. Montana
Department of Biomedical Engineering,
King's College London,
St Thomas' Hospital, London SE1 7EH, UK
E-mail: giovanni.montana@kcl.ac.uk